

Hypermedia-Aided Design

Darko Kirovski, Milenko Drinić, and Miodrag Potkonjak
Computer Science Department
University of California, Los Angeles, CA 90095

ABSTRACT

Recently, the Internet revolutionized many activities from entertainment to marketing and business. Two key underlying Internet technologies, efficient data delivery and hypertext, demonstrated exceptional potential as new application enablers. In this paper, we present a novel Hypermedia-Aided Design (HAD) collaboration framework that facilitates new communication and data presentation paradigms to improve the effectiveness of typical EDA applications. The framework leverages on the advantages of using semantic multicast as a communication backbone and quantized hypermedia presentations as an efficient data organization, retrieval, and presentation model. Semantic multicast is a global communication tool that relies on an inter-network of proxies to provide content discovery and semantics-based profile-driven data dissemination services. We introduce the notion of a quant, an atomic interactive multimedia information primitive with embedded hyperlinks. We demonstrate how interest-specific quant retrieval and concatenation can enable more focused collaboration.

The HAD framework consists of a set of applications for student(designer)-centered CAD education(consulting), collaborative design and debugging, I-commerce, and technical support. In all applications, quant-based presentations enable that theoretical and practical components are tailored according to user's interest and performances. The conceptualized and implemented applications act in synergy with existing software, hardware, and system design tools.

1. INTRODUCTION

Ever since the occurrence of first written documents, the perceptual and organizational structure of the collected human knowledge has corresponded to the social and technical maturity of the human civilization. Evolution of photography, motion picture, television, video, and interactive systems has stimulated revolution in many areas of human activity. The latest emergence of the Internet as the ultimate communication paradigm and source of information, and hypermedia as the most efficient way to organize and present media, has provided an engine for development of applications of the new era.

Due to the emerging semiconductor technologies and shortened time-to-market, the new era for the EDA industry is determined towards facilitating design reuse and balancing the implementation trade-offs among technologies such as programmable and application-specific cores. System integration strategies have affected the business model of virtually all EDA and

semiconductor companies [VSI98]. For example, recently, a number of companies have consolidated their efforts towards developing off-the-shelf programmable cores (e.g. ARM, LSI Logic). Significant research attention has been paid to enabling efficient integration of such systems [Jha95, Kis97, Gup97] and providing a secure intellectual property (IP) business model [VSI98, Kah98, Hon98]. However, only little effort has been focused towards determining the communication environment for a design process that depends on IP blocks and tools [Spi97, Sah97, Hin98].

In this paper, we present the concepts and implementation of the key enabling technologies for effective research, development, and technical support of IP-based EDA companies. We have adopted a novel communication paradigm, semantic multicast, as a solution to the content discovery, data dissemination, and security problems. In opposition to multicast [Obr98], here the transmitted data is analyzed at an inter-network of proxies which provide the dissemination and announcement services.

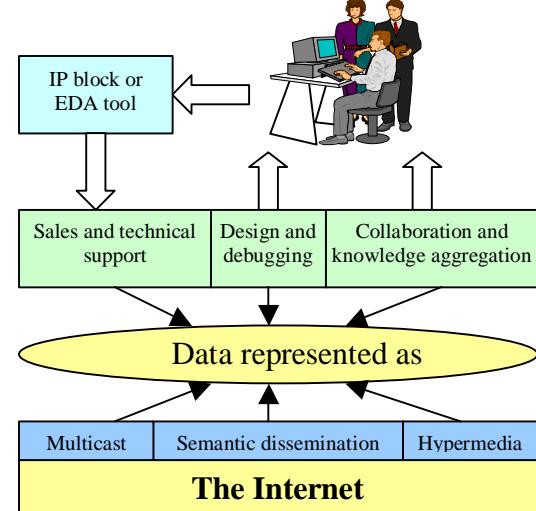


Figure 1. The enabling data communications, dissemination, and presentation framework for the EDA IP business model.

We introduce the notion of a quant, an atomic on-line multimedia presentation primitive with embedded hyperlinks, and show how it can exploit the advantages of the Internet and hypermedia while providing support for efficient business operability of an IP-based EDA or semiconductor company. The quants enable context reuse, selectivity, and customization at the both the data retrieval and presentation level. In addition, data presentation systems that facilitate quants enable effective statistical evaluation for extraction of user profiles, design- and tool-specific usage metrics. We demonstrate the rules for effective quant recording, organization, retrieval, and presentation.

We have developed a set of applications that facilitate the novel underlying communication framework and data presentation model. The applications provide the following services:

- **just-in-time education (consulting),**
- **hypermedia-aided event-driven system design,** and
- **Internet-commerce and technical support.**

Using this set of applications we have conducted a course on intellectual property protection techniques, enabled an effective distance-consulting tool, and created an interactive visual wrapper for a design-watermarking tool.

2. RELATED WORK

Collaboration in the EDA environment has been recognized as one of the key enabling technologies for short time-to-market. Spiller and Newton have identified a number of key research and development challenges in employing the Internet in collaborative EDA [Spi97]. A number of WWW-based design systems have been introduced: WELD [Cha98], the distributed microsystem design [Sah97, Mic98] and validation [Hin98] projects. Benini et al. have demonstrated an interface for on-line integration of EDA tools [Ben96]. They have focused more on interactive remote execution than on information retrieval. Marketing IP products on-line is already a reality. For example, The Design and Reuse Co. provides an on-line IP catalog, where IP blocks can be evaluated, purchased, and exchanged [Des98].

The impact of effective technical support on product marketability has been addressed by Wilson [Wil94]. He presented data on how organizing and automating the process of keeping track of problems, their solutions, and rationalizing their dissemination may pay big dividends.

3. UNDERLYING ENABLING TECHNOLOGIES

3.1 Semantic Multicast

The paradigm of multicasting information on the Internet has been supported by a number of protocols [Obr98]. We have conceptualized and implemented semantic multicast as a means to realize a large-scale shared interaction infrastructure for seamless collection, indexing, and dissemination of data produced in multicast collaborative sessions. A collaborative session is considered to consist of two collections of objects:

- a set of users with: interests, degree of interest, and abilities to participate in the streams of a collaborative session and
- a set of interaction streams arising in the session that needs to be disseminated to users in various forms.

The goal of semantic multicast communication framework (implemented as a network of intermediate proxies) is to create a logical dissemination, filtering, and archiving structure for making streams of collaborative sessions available to correct users at the right amount of detail and in an efficient manner. Given a collaborative session that consists of many streams on a single multicast channel, the semantic multicast inter-network of proxies establishes a semantic flow graph of collaborative streams between interest groups. Users are subscribed to appropriate proxies based on their profiles. The collaborative session becomes a multi-level multicast of data from sources through proxies and to user interest groups. The proxy archives each stream; filters it for a specific semantic topic, merges it with similar streams, and disseminates it to all users subscribed to the proxy. As the proxy archives the stream, it performs a more detailed off-line analysis to provide additional semantic structuring for subsequent retrieval and feedback to the semantic multicast graph. Details about the associated enhancements to the multicast protocol and real-time audio, video, and text filtering algo-

rithms can be found in [Dao98]. Most importantly, the semantic multicast communication framework satisfies the five fundamental criteria (realized in [Spi97]) for effective application to the EDA industry: scalability, availability, adaptability, robustness, and cost effectiveness.

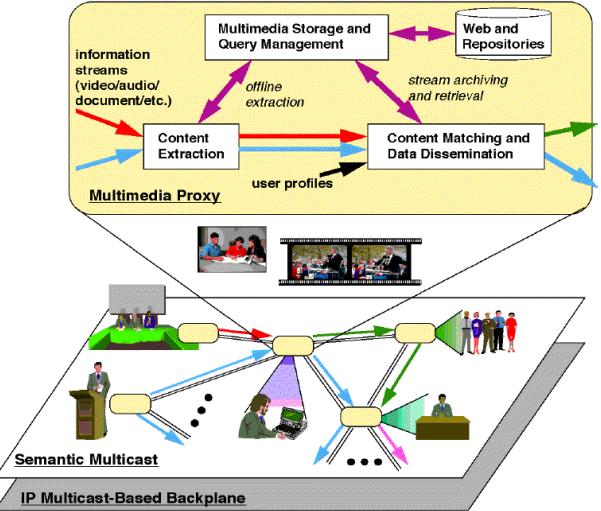


Figure 2. Information dissemination in semantic multicast.

Figure 2 depicts the general components comprising the semantic multicast process. Specifically, an information stream enters the network at any node. Each stream is subject to a content extraction process to approximate its general semantics. Based on the content summary, the stream is matched against the semantic coverage of proxies, annotated, and multicast to the appropriate proxy in the network. Upon receiving the annotated information stream, the proxy archives the stream and filters it into a level of detail appropriate for dissemination to its user groups. A proxy may choose to merge two or more incoming information streams and disseminate (and archive) the fused data product. The proxy archives and performs detailed offline content analysis and indexing of the stream. As new detailed semantics are uncovered in the stream, this information is fed back into the content matching process to see if additional proxies are "interested" in the discovered content. At any time, a user may interact with a proxy and request specialized stream summaries to be created and sent to the user.

3.2 Quant-based Presentation

Hyperlinks in text (hypertext) and multimedia (hypermedia) have been already demonstrated to qualitatively and quantitatively improve the organization and expressiveness of data [Har94, Hal94]. In order to facilitate the advantages of using multimedia and hyperlinks in presenting information, we introduce a novel data presentation and organization model: quants.

A **quant** is an atomic on-line multimedia presentation primitive with embedded hyperlinks that lasts finitely in time and conveys an atomic knowledge instance. In more detail, a quant can be defined using the following set of rules:

- **Continuity.** A quant contains one or a synchronized continuous set of time lasting media (video/audio). Therefore, a quant has the property of duration in time.
- **Synchronicity.** Other non-time-lasting media used to convey information is synchronized with the time lasting actor. An

example of a set of media that can be supported for quant creation is presented in Figure 3. As shown, for example, at time $t=5\text{sec}$ a slide appears in the slide viewer or at time $t=34\text{sec}$ an event is spawned in the Java applet.

- **Global reusability.** The media files contained in a quant are accessible on-line using URL pointers. This enables global reuse of media files.
- **Monolithic viewing.** The user may be provided with all tools for viewing time lasting media (fast forward, reverse viewing, etc.). The media encompassed in a quant is monolithic with respect to time; i.e. viewing the quant in fast forward or reverse mode should be accompanied with synchronized appearance of all non-time-lasting media.
- **Interactivity.** A presentation of a single quant may be ended with an opportunity to interact with the server (for example using a Java applet that enables form filling, messaging, or Internet telephony).
- **Atomicity.** A quant should convey an atomic instance of knowledge. Quants are envisioned to solve the problem of bulkiness of audio and video content by quantization into shorter independent units. For example, a lesson may consist of several quants, each describing completely a particular sub-topic within the subject.

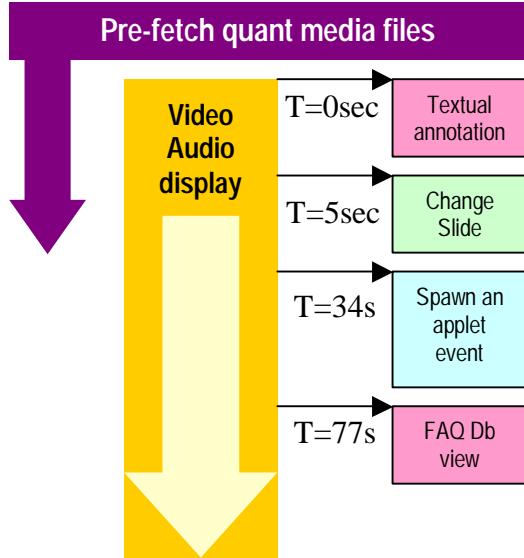


Figure 3. An example of a quant. The quant viewing procedure consists of three processes: pre-fetching the media files, audio/video play, and synchronized display of non-time-lasting media.

A quant-based presentation can be realized, for example, as a list, tree, or Petri net structure of quants offered to a client. The client (user or his recommendation system, e.g. the semantically filtered multicast engine) selects a subset from the list of quants for viewing. The client is provided with the ability to temporarily stop the quant viewing and annotate that moment in the presentation for future reference. Therefore, the user maintains a local database of multimedia annotations for a particular quant.

The multimedia information within a quant can be augmented with hyperlinks, which point to other multimedia data, quants, or quant presentations. Upon following a hyperlink towards another quant or a presentation, the current quant is stopped and the new quant or the first quant of the linked presentation is loaded and viewed. The quants are stacked in

entation is loaded and viewed. The quants are stacked in such a way that upon termination of the current quant, the control returns to the next quant in the stack in the moment when it was interrupted.

An important qualification of the effectiveness of our implementation of quant-based presentations can be traced along the seven design concepts for navigation in cyberspace [Glo97]:

- **Linking.** Due to content quantization, quants enable effective structuring of multimedia content using both static and dynamic [Fou90] links. Links are embedded into presentations to: (a) itemize the list of quants or groups of quants, (b) hierarchically present the encompassed set of quants, (c) point from one quant media to another quant with similar content, and (d) express the results of a search for quants on a particular topic.

- **Searching.** The traditional search engines [Har98] can be used for content discovery in a pool of quants. In addition, we employ sophisticated video scene change and approximate SQL queries for text filtering to provide search capabilities adequate to multimedia documents [Dao98].

- **Sequentialization.** According to a given user-profile, his or her background and interest, the system assembles a guided tour [Hal88] of a given presentation. The user is able to modify the tour upon request.

- **Hierarchy.** We have adopted the multi-tree paradigm [Fur94] for constructing hierarchically a presentation from a pool of quants. Each tree is built with the intent to support one guided tour.

- **Similarity.** The approximate queries supported by our search engine explore similarities among a pair of quants as well as a quant and a user profile. While in the latter case the similarity functions are determined by the user, in the former case, we used the electronic index approach [Sal88]. In both cases the similarity $\text{sim}(I,J)$ between two objects I and J has to be described in such a way that:

1. $\text{sim}(i,i) \geq (\forall j)\text{sim}(i,j)$
2. $\text{sim}(i,j) = \text{sim}(j,i)$
3. $(\forall i,j)\text{sim}(i,j) \geq 0$

We define the inverse gravitation between two strings as a number of matched words in both strings divided by their length and use it as a heuristic similarity function:

$$\text{sim}(i,j) = G(i,j) = \frac{\text{words_matched}^2(i,j)}{\text{len}(i) \cdot \text{len}(j)}$$

- **Mapping.** Currently, the system for fish-eye visualization [Sar92] of quant based presentations is under development. We intend to compute the degree of interest **DOI()** function [Joh87] that guides the sizing and positioning of presented objects on the screen using parameters indicated in the user profiles and quant definition file.

- **Agents.** Currently, we support only an autonomous agent which stores information broadcast over a set of multicast channels that is of user's interest [Dao98]. The information can be retrieved at user request. Note that this feature is partially enabled by the semantic multicast proxies (content dissemination and advertising) and partially by a user-local storage system [Dao98].

4. HYPERMEDIA-AIDED APPLICATIONS for EDA

In this section, based on the communications and data organization and presentation primitives discussed in Section 3, we present the technical details of the implemented applications for

student(designer)-centric distance learning(consulting), hypermedia-aided system design and debugging, and EDA product tech support.

4.1 Student(Designer)-Centric Distance Learning

The research and development activity in the EDA industry generates information of exceptional volume and diversity. Current push-pull Internet technologies are not addressing the problem of data collection nor dissemination well. Similarly, modern search engines provide a potentially extensive view on the available information, but fail to support just-in-time retrieval and ease of usage.

Using the semantic multicast communication framework, we have conceptualized a real-time quant-based information discovery, presentation, and collaboration application. The target

according to his or her personal interest profile, the offered multi-tree of quants and associated data may be modified (filtered) seamlessly by the network proxy agents with respect to the extracted content of the broadcast media. Once the modified list of quants is received, the student is able to select a subset of quants for viewing. The professor is able to monitor the selection of quants for each student. The student is provided with the ability to stop viewing a quant at any moment, and add a time-stamped multimedia annotation. This annotation is kept local and the students are encouraged to reduce the security locks of annotations so that the professor may have better overview of their understanding of the material.

The interaction between the professor and the students is enabled using a set of real-time multicast communication tools: video conferencing for high-speed or Java telephony [Jav98t]

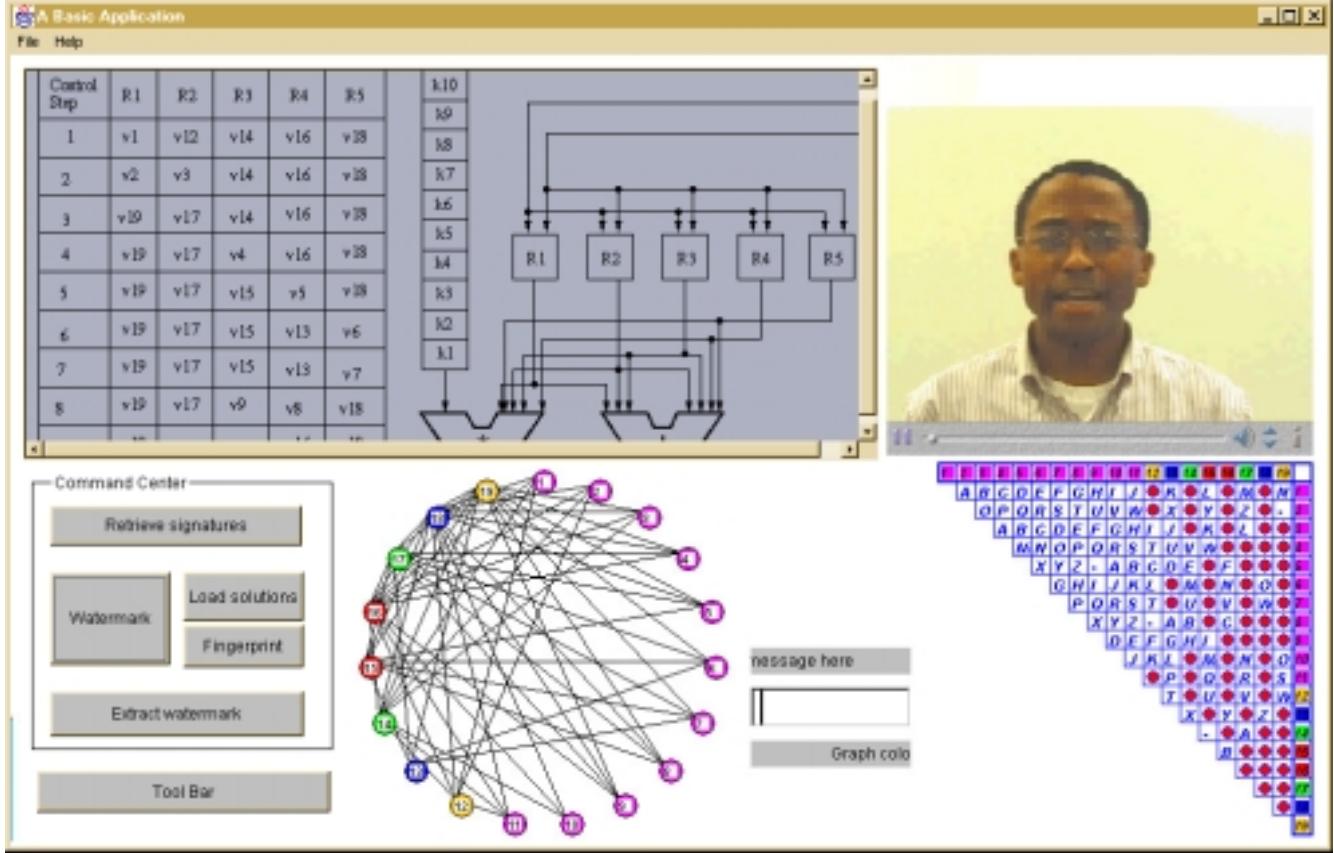


Figure 4. QuantViewer: presenting a quant that explains the idea behind IPP using watermarking and provides a shell for a shareware graph coloring watermarking software.

application for distance learning supports the following collaboration model. Two objects can participate in a learning session: a student and a professor. The session consists of two independent intercommunication fragments: quant-based presentation and student-professor interaction. Initially, the professor prepares a list of quants and creates a multicast session for a presentation on a particular topic. Next, the professor announces the session on an announcement multicast channel according to the multicast protocol [Obr98]. A number of interested students subscribe to the session. While the professor is responsible for uploading the quant data into the session channel, the student is required to store the quant data locally for future display. For each student,

and messaging [Jav98m] for low-speed communications, whiteboard for still image presentation and discussion [Li98], etc. We used the COCA virtual machine and its collaboration specification language to control the collaboration policies among participants [Li98]. The COCA virtual machine acts as a middleware framework between the interactive collaborative applications and the communication sockets. The collaboration model is specified in the COCA language, which is at run-time interpreted to create a desired set of collaboration policies. Such approach enables secure and effective granting and restricting access to information during the collaboration.

The current implementation of this collaboration framework encompasses a quant viewing application and an interactive COCA-controlled audio-whiteboard collaboration tool. The quant viewer is specifically tuned for the EDA tasks. It enables review of synchronized video, audio, slides with embedded hyperlinks, hypertext, FAQ database views, and, most importantly Java applets. The applets can be used to present simulations, EDA tool GUIs (the EDA tool can be written in C/C++), and enable form filling.

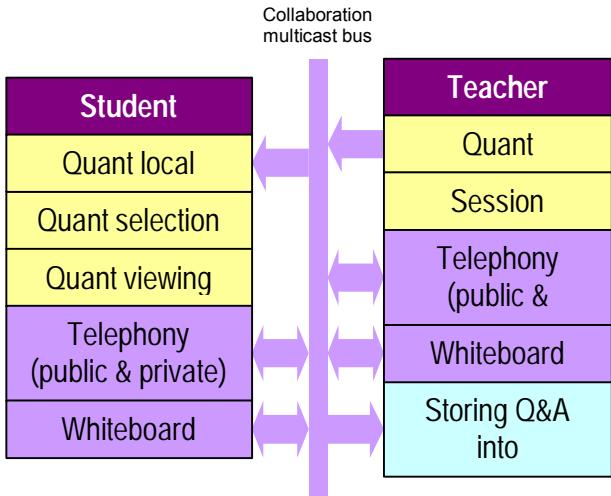


Figure 5. Control flow diagram of the implemented prototype for the remote consulting application.

We have conducted a *course on techniques for intellectual property protection* primarily based on a set of recorded quants. Similarly, we have used this set of quants to create a *wrapper around a design watermarking tool prototype*. The quant viewer in this case creates a working environment for guidance, help, and execution for watermarking of behavioral synthesis solutions. In addition, an on-line learning session can be enabled using an application, which facilitates the standard MBone [Obr98] and/or the developed semantic multicast communication framework. The available tools for interactive communication include Java telephony [Jav98t] and a whiteboard. The interaction of the professor and student roles with respect to the collaboration events is illustrated in Figure 5. The application for viewing quants is visualized using Figure 4; while displaying a quant which encapsulates an audio/video presentation accompanied with a synchronized display of a set of slides, a demo applet, and a shell for actual application of a graph coloring watermarking tool.

We have augmented the distance learning application with available electronic commerce tools [Jav98c] to support *consulting services*. Using the demonstrated set of tools, we have enriched the existing consulting and licensing business model [Bet98] to the pay-per-minute and pay-per-quant licensing scheme. Measuring the duration of the collaboration supports the first consulting model. Due to the rich content of the information contained in a quant, we have found that the second consulting model can be also effective in many circumstances. Finally, we have developed a novel quant-based system for *on-line sales and tech support* that utilizes customer profiles, inquiry complexity evaluation, intelligent FAQ database, and quant demos to provide inexpensive semi-automated tech support system.

4.2 Hypermedia-aided Design and Debugging

Modern designs are converging towards the concept of a core-based programmable system-on-chip. Designing and debugging such systems is an exceptionally difficult and time-consuming task. We have conceptualized a hypermedia-aided design and debugging tool, which enables the user to have an event-centric view on the developed architecture at various design levels using a spectrum of hyperlinks. In addition, we have developed a tool for on-line documentation maintenance, which employs multimedia and hyperlinks rather than only textual annotations.

4.2.1 Hyperlinks in EDA tools and designs

Modern IP blocks are usually marketed as VHDL black-box components accompanied with C-model simulators for fast functional execution during debugging (e.g. Escalade's IP Guard [Esc98]). The IP developer is responsible for providing an extensive off-line documentation, which is most often cumbersome and rarely meets the system integrator requirements. We have conceptualized an application for IP block design, which facilitates hyperlinks to hypermedia documents to provide better understanding of the design architecture and its response to external events. The goal is to enhance the design environment in a way that required information is always available on-line, easily understandable, and recorded using the most effective media (text, image, voice, animation, etc.). We identify two different classes of hyperlinks associated to design objects:

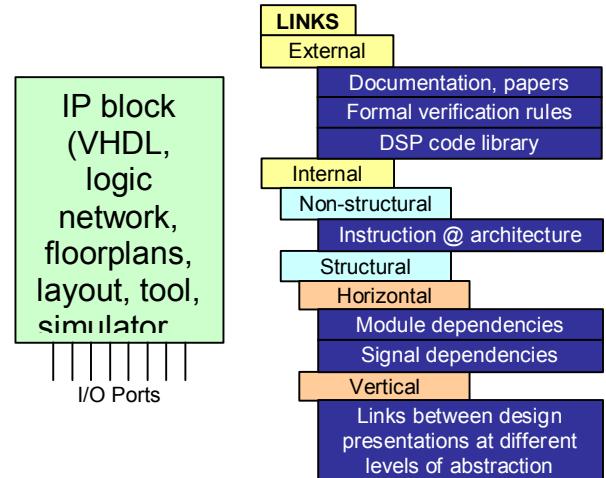


Figure 6. An IP block and its associated set of hyperlinks to hypermedia artifacts.

- **External hyperlinks** are associated to objects that are assigned artifacts that describe some of their properties, external with respect to the IP block. An example of such a hyperlink is a pointer to the complete user guide, code library, etc. External hyperlinks are by default readable by the world.
- **Internal hyperlinks** are associated with a particular part of the design, a dependency between two parts of the design, or a dependency between a part of the design and an external signal. Commonly the IP developer does not want to reveal all internal links to the customer. Internal hyperlinks are further categorized into:
 - **Non-structural hyperlinks** are possibly related to the entire or a large part of the design. An example of such a hyperlink would be a link from an instruction in the assembly editor to a marked architecture schematic, HDL code, logic network, simulator C code, and/or list of hazards and stalls which is

closely related to this instruction. At debugging time, useful hyperlinks are ones associated to external events (interrupts, control signals) which incur specific behavior of usually only a part of the design. For example, an interrupt signal should be hyperlinked to the interrupt service routine which services it and modified locations in the memory map (buffer loaded). Links among assembly or higher-level language procedures are also desired for better overview of program-architecture interdependencies.

- **Structural hyperlinks** are useful at debugging time while browsing the design's architecture at various levels of design abstraction. Hyperlinks between modules at the same level of design abstraction are called **horizontal**, while hyperlinks between a higher and a specific lower-level design construct are called **vertical**. Structural hyperlinks are useful in pinpointing to critical building blocks of a higher level module (e.g. critical path), describing protocols (e.g. several animations may describe more effectively all scenarios of a communication protocol than a textual explanation).

In general, all links should assist the IP block developer and the system integrator to conveniently browse portions of the design and maintain a collection of documents about the properties of the IP block, examples of usage, and module and signal interdependencies. Using hyperlinks and multimedia, the documentation should be made available during design and integration in a format, which conveys the knowledge in the most effective manner. An example of a set of hyperlinks associated with a typical IP block is presented in Figure 6.

5. CONCLUSION

Although the state of current semiconductor technologies provides a mature foundation for application of the design-and-reuse system development paradigm, many other aspects of such a development process are still in an early development stage. In this paper, we have introduced a vision of a set of applications, which facilitate a semantically filtered multicast backbone and a novel hypermedia presentation model along, quants which effectively enable the crucial EDA tasks: collaboration, design and debugging, documentation maintenance, e-commerce, and technical support. To show the advantages of the new data presentation model, we have implemented in Java a subset of the conceptualized support applications.

REFERENCES

- [Ben96] L. Benini, A. Bogliolo, and G. De Micheli. Distributed EDA tool integration: the PPP paradigm. International Conference on Computer Design, pp.448-53, 1996.
- [Bet98] D.R. Bettadpur. Software Licensing Models in the EDA Industry. South Asia and Pacific Design Automation Conference, 1998.
- [Cha98] F.L. Chan, M.D. Spiller, and A.R. Newton. WELD - An Environment for Distributed Electronic Design. Design Automation Conference, pp.146-51, 1998. <http://www-cad.EECS.Berkeley.EDU/Respep/Research/weld/>
- [Dao98] S. Dao et al. Semantic Multicast - Intelligently Sharing Collaborative Sessions. 1998. http://www.wins.hrl.com/projects/semcast/semcast_whitepaper.ps
- [Des98] <http://www.design-reuse.com/>
- [Esc98] <http://www.escalade.com/>
- [Fou90] A.M. Fountain, W. Hall, I. Heath, and H.C. Davis. MICROCOSM: an open model for hypermedia with dynamic linking. Hypertext: Concepts, Systems and Applications. European Conference on Hypertext, pp.298-311, 1990.
- [Fur94] G.W. Furnas and J. Zacks. Multitrees: enriching and reusing hierarchical structure. CHI '94 Conference Proceedings, pp.330-6, 1994.
- [Glo97] P. Gloo. Elements of hypermedia design: techniques for navigation & visualization in cyberspace. Boston, Birkhäuser, 1997.
- [Gup97] R.K. Gupta and Y. Zorian. Introducing core-based system design. IEEE Design & Test of Computers, Vol.14, (no.4), pp.15-25, 1997.
- [Hal88] F.G. Halasz. Reflections on NoteCards: seven issues for the next generation of hypermedia systems. Communications of the ACM, Vol.31, (no.7), pp.836-52, 1988.
- [Hal94] F. Halasz and M. Schwartz. The Dexter Hypertext Reference Model. Communications of the ACM, Vol.37, (no.2), pp.30-9, 1994.
- [Har94] L. Hardman, D.C. Bulterman, and G. van Rossum. The Amsterdam hypermedia model: adding time and context to the Dexter model. Communications of the ACM, Vol.37, (no.2), pp.50-62, 1994.
<http://harvest.cs.colorado.edu/>
- [Hon98] I. Hong and M. Potkonjak. Intellectual Property Protection Techniques for Behavioral Specifications. Unpublished manuscript, 1998.
- [Jav98c] <http://java.sun.com:80/products/commerce/docs/index.html>
- [Jav98m] <http://java.sun.com/products/jms/docs.html>
- [Jav98t] <http://java.sun.com:81/products/jtapi/download2.html>
- [Jha95] P.K. Jha and N.D. Dutt. Design reuse through high-level library mapping. European Design and Test Conference, pp.345-50, 1995.
- [Joh87] W.P. Jones and G.W. Furnas. Pictures of relevance: a geometric analysis of similarity measures. Journal of the American Society for Information Science, Vol.38, (no.6), pp.420-42, 1987.
- [Kah98] A.B. Kahng et al. Robust IP Watermarking Methodologies or Physical Design. Design Automation Conference, 1998.
- [Kis97] P. Kission and A.A. Jerraya. Behavioral design allowing modularity and component reuse. Journal of Microelectronic Systems Integration, Vol.5, (no.2), pp.67-83, 1997.
- [Li98] D. Li and R.R. Muntz. COCA: Collaborative Objects Coordination Architecture. To appear ACM CSCW, 1998.
- [Mic98] http://web_nt.sainc.com/arpa/msdproject/
- [Obr98] K. Obraczka. Multicast transport protocols: a survey and taxonomy. IEEE Communications Magazine, Vol.36, (no.1), pp.94-102, 1998.
- [Sah97] D. Saha and A.P. Chandrakasan. A Framework for Distributed Web-based Microsystem Design, Sixth Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises, 1997.
- [Sal88] G. Salton. Automatic text processing: the transformation, analysis, retrieval of information by computer. Addison-Wesley, 1988.
- [Sar92] M. Sarkar and M.H. Graphical fisheye views of graphs. CHI '92 Conference Proceedings, pp.83-91, 1992.
- [Spi97] M.D. Spiller and A.R. Newton. EDA and the Network, Design Automation Conference, pp.470-5, 1998.
- [VSI98] <http://www.vsi.org>
- [Wil94] R. Wilson. The fine art of technical support. CD-ROM Professional, Vol.7, (no.1), pp.91-2, 1994.